

(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 1 164 746 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
19.12.2001 Bulletin 2001/51

(51) Int Cl.7: **H04L 9/32**(21) Application number: **01119418.0**(22) Date of filing: **01.11.1996**

(84) Designated Contracting States:
**AT BE CH DE DK ES FI FR GB GR IE IT LI LU MC
NL PT SE**

(72) Inventor: **Micali, Silvio**
Brookline, MA 02146 (US)

(30) Priority: **02.11.1995 US 6143 P**
11.10.1996 US 729619

(74) Representative: **Bickel, Michael**
Westphal - Mussnug & Partner Patentanwälte
Mozartstrasse 8
80336 München (DE)

(62) Document number(s) of the earlier application(s) in
accordance with Art. 76 EPC:
96937813.2 / 0 858 702

Remarks:

This application was filed on 13 - 08 - 2001 as a
divisional application to the application mentioned
under INID code 62.

(71) Applicant: **Micali, Silvio**
Brookline, MA 02146 (US)

(54) Tree-based certificate revocation system

(57) A method for a intermediary to prove to an end
user that certificate information is authenticated by an
authority, comprising the steps of:

(a) obtaining at least a portion of an authenticated
tree having certificate nodes corresponding to cer-

tificate values indicative of the information; and
(b) causing the end user to receive at least one of
the following values: certificate values, authenticat-
ing values of certificate values, and one or more
node values authenticated by an authority.

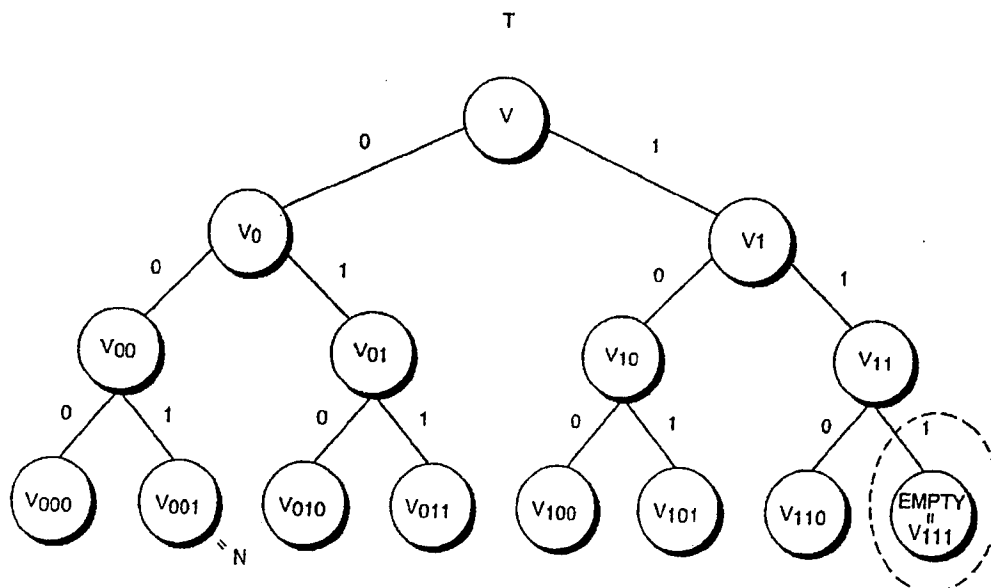


FIG. 1

EP 1 164 746 A2

Description

[0001] This application is based on U.S. provisional patent application No. 60/006,143 filed on November 2, 1995.

TECHNICAL FIELD

[0002] The present invention relates generally to secure communications and more particularly to schemes for certificate management.

BACKGROUND OF THE INVENTION

[0003] In many settings, it is useful to certify data, as well as to revoke data that was previously certified. For instance, in a Public Key Infrastructure (PKI), it may be useful to certify users' public keys. Such certification may be provided in the form of a certificate which contains the certified data and vouches authenticity of the certified data.

[0004] In a digital signature scheme, each user U chooses a signing key SK_U and a matching verification key, PK_U . User U uses SK_U to compute a digital signature of a message m , $SIG_U(m)$, while anyone knowing that PK_U is U's public key can verify that $SIG_U(m)$ is U's signature of m . Finding $SIG_U(m)$ without knowing SK_U is practically impossible. On the other hand, knowledge of PK_U does not give any practical advantage in computing SK_U . For this reason, it is in U's interest to keep SK_U secret (so that only he can digitally sign for U) and to make PK_U as public as possible (so that everyone dealing with U can verify U's digital signatures). At the same time, in a world with millions of users, it is essential in the smooth flow of business and communications to be certain that PK_U really is the legitimate key of user U. To this end, users' public keys are often "certified" by a certificate that serves as proof that U is the legitimate owner of PK_U . At the same time it is also useful to be able to revoke some of the already-issued certificates when U is no longer the legitimate owner of PK_U (for whatever reason) and/or when SK_U has been compromised. Of course, the need for certification and certificate revocation extends beyond certifying public keys.

[0005] In many instances, certificates for users' public keys are produced and revoked by certifying authorities called CA's. A complete public key infrastructure may involved other authorities (e.g., PCAs) who may also provide similar services (e.g., they may certify the public keys of their CA's). The present discussion can be easily applied to such other authorities in a straightforward manner.

[0006] A CA may be a trusted agent having an already certified (or universally known) public key. To certify that PK_U is U's public key, a CA typically digitally signs PK_U together with (e.g., concatenating it with) U's name, a certificate serial number, the current date (i.e., the certification or issue date), and an expiration date. The CA's

signature of PK_U is then inserted in a Directory and/or given to U himself. Note that, before certifying U's public key, it is necessary to perform additional steps, such as properly identifying user U. However, these additional steps are optional.

[0007] Upon receiving the (alleged) digital signature of user U of a message M , $SIG_U(M)$, a recipient R needs to obtain a certificate for PK_U . (In fact, $SIG_U(M)$ may be a correct digital signature of M with respect to some public key PK_U , but R has no guarantee that PK_U is indeed U's public key. Recipient R may obtain this certificate from the Directory, or from his own memory (if he has previously cached it), or from U himself. Having done this, R verifies (1) the correctness of the CA's certificate for PK_U with respect to the CA's public key, and (2) the correctness of $SIG_U(M)$ with respect to PK_U . If the CA's public key is not universally known, or cached with R, then a certificate for the CA's key may also be obtained.

[0008] Certificate retrieval is thus possible, although not necessarily cheap. Unfortunately, however, this is not the only retrieval that R needs to do. In addition, it is important that R makes sure that the certificate for PK_U has not been revoked. This check, of course, may not be needed after the certificate's expiration date, but may be needed during the certificate's alleged lifetime. A user's certificate can be revoked for a variety of reasons, including key compromise and the fact that the user is no longer associated with a particular CA.

[0009] To enable a recipient to establish whether a given certificate has been revoked, it is known to have each CA periodically issues a Certificate Revocation List (CRL for short). A CRL may consist of the issuer's digital signature of a header comprising the issuer's name (as well as the type of his signature algorithm), the current date, the date of the last update, and the date of the next update, together with a complete list of revoked certificates (whose date has not yet expired), each with its serial number and revocation date. Since it is expected that a CA revokes many certificates, a CRL is expected to be quite long. It is envisaged that the CRL is provided to a directory who may then distribute the CRL to end users.

[0010] After performing some checks on the CA's CRL (e.g., checking the CA's digital signature, checking that the CRL has arrived at the expected time, that a certificate declared revoked in the previous CRL of that CA - and not yet expired - still is revoked in the current CRL, etc.), the Directory stores it under its CA name.

[0011] When a user queries the Directory about the revocation of a certificate issued by a given CA, the Directory responds by sending to the user the latest CRL of that CA. The user can then check the CRL signature, the CRL dates (so as to receive a reasonable assurance that he is dealing with the latest one), and whether or not the certificate of interest to him belongs to it.

[0012] While CRLs are quite effective in helping users establishing which certificates are no longer deemed valid, they are also extremely expensive, because they

tend to be very long and need to be transmitted very often.

[0013] The National Institute of Standard and Technology has tasked the MITRE Corporation to study the organization and cost of a Public Key Infrastructure (PKI) for the Federal Government. This study estimates that CRLs constitute by far the largest entry in the Federal PKI's cost list. According to MITRE's estimates/assumptions, in the Federal PKI there are about three million users, each CA serves 30,000 users, 10% of the certificates are revoked (5% because of key compromise and 5% because of change in affiliation with the organization connected to a given CA), CRLs are sent out bi-weekly, and the recipient of a digital signature requests certificate information 20% of the time (assuming that the remaining 80% of the time he will be dealing with public keys in his cache). The study envisages that each revoked certificate is specified in a CRL by means of about 9 bytes: 20 bits of serial number and 48 bits of revocation date. Thus, in the Federal PKI, each CRL is expected to comprise thousands of certificate serial numbers and their revocation dates; the header, however, has a fixed length, consisting of just 51 bytes.

[0014] At two cents per kilobyte, the impact of CRL transmission on the estimated yearly costs of running the Federal PKI is stunning: if each federal employee verifies one hundred digital signatures per day on average, then the total PKI yearly costs are \$10,848 million of which 10,237 million is due to CRL transmission. If each employee is assumed to verify just five digital signatures a day on average, then the total PKI yearly costs are \$732 million, of which 563 million is due to CRL transmission.

[0015] The MITRE study thus suggests that any effort should be made to find designs alternative to and cheaper than conventional CRL's.

[0016] In addition, we contend that it is possible for a user to query the Directory with a serial number not corresponding to any issued certificate. (Indeed, while many times the user has already seen a certificate and accesses the Directory just to confirm the current validity of that certificate, at other times the user wishes to obtain the corresponding certificate from the Directory). If the corresponding certificate does not exist, the Directory is at a loss as to how to proceed. If the Directory responds truthfully, it may not be believed by the user. If the Directory gives the users all the certificates in its possession (or those relative to a given CA) the user may suspect that the Directory left out the certificate of interest. Indeed, even if the Directory gives the user the latest CRL of a given CA, this does not prove to the user that the certificate in question does not exist. (In fact, the actions of the Directory may actually be interpreted as saying that the certificate is valid because it does not appear to have been revoked.) Thus in this thorny situation the Directory would have to be trusted.

BRIEF SUMMARY OF THE INVENTION

[0017] To avoid the dramatic CRL costs, a novel *Certification Revocation System* is described, where requesting users no longer receive the latest list of revoked certificates (of a given CA). The scheme utilizes a known tree-based authentication technique in a novel way to overcome the problems associated with the prior art.

[0018] It is thus a primary object of this invention to provide certificate management without providing CRL's to a user requesting information about the certificate (e.g., its validity). Although special CRL's still may be used between CA's and the Directory in this scheme, the tree-based technique allows the Directory to convince users of whether or not a given certificate is still valid in a way that is essentially individualized, and thus quite compact and convenient.

BRIEF DESCRIPTION OF DRAWINGS

[0019] The sole figure shows a Merkle tree that is used in connection with the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0020] An interesting authentication scheme is described in U.S. Patent No. 4,309,569 to Merkle where it is used in connection with digital signatures. The sole figure shows a simple Merkle tree, T. Merkle is hereby incorporated by reference and familiarity with the Merkle scheme is assumed in the following discussion. The Merkle scheme has been shown to be useful for applications other than digital signature schemes; in particular for constructing special and more efficient mathematical proof systems. The object of the present invention is to show how Merkle's scheme can yield certificate revocation systems more efficient than known CRL-based systems.

[0021] Briefly, a Merkle tree consists of a tree (binary, ternary, binary and ternary, etc.) with values corresponding to the tree nodes in the following manner. For simplicity, consider a full binary tree T with $n = 2^k$ leaves and let H be a one-way hash function mapping strings of arbitrary length into B-bit strings (that is, a function H for which it is difficult to find two different strings x and y such that $H(x) = H(y)$). Then, T is a Merkle tree if the value corresponding to each internal node equals the one-way hash of the values corresponding to the children thereof. Developing a minimum of terminology, let N be an internal node where the left child is L and the right child is R, and let V_L be the value corresponding to L and V_R the value corresponding to R (we shall refer to V_L and V_R as, respectively, a left-value and a right-value). Then, for T to be a Merkle tree, the value corresponding to node N must be the B-bit value $H(V_L V_R)$, where $V_L V_R$ is the concatenation of V_L and V_R (although

V_L and V_R may be combined by operations other than concatenation).

[0022] It is a well known property of a Merkle tree that (unless one "breaks" the one-way hash function -- i.e., unless one is capable of finding two different strings x and y such that $H(x) = H(y)$ -- which would require an extraordinary amount of computation to do) changing the value of any node in a Merkle tree causes the root value to change also. It is also known that to verify whether the value of a given node N is authentic with respect to the root value of a Merkle tree, it suffices to use very few values of the tree. Indeed, it suffices to use the *authentication path* of node N , which is the sequence of values corresponding to the siblings of the nodes along the path from node N to the root. Each of the values of the nodes in the authentication path is an *authentication value*. If most of the internal nodes of a Merkle tree have two children, it is immediately seen that an authentication path comprises roughly k authentication values where k is the depth of the tree, even though the total number of nodes can be as much as 2^k . (To facilitate the verification of the value corresponding to a node N whose position within the tree is unknown, the sequence of authentication values for N can be given by specifying whether each value is a left value or a right value. Vice versa, it is well known that if the authentication path of node N is verified as valid, the fact that each authentication value is a left or right value determines the position of node N within the tree.)

[0023] Let us now recall, by way of example, how one can verify that the value of a node N is correct with respect to the value of the root. Referring to the sole figure, let N be the second leaf from the left in the tree T . Then, V_{001} is the value of N ; V is the value of the root; and the authentication path of N consists of the following sequence of three values: the left value V_{000} , the right value V_{01} , and the right value V_1 . Then, because V_{000} is a left value, one computes $H(V_{000}, V_{001})$ and calls V_{00} the result. Then, because V_{01} is a right value, one computes $H(V_{00}, V_{01})$ and calls the result V_0 . Finally, because V_1 is a right value, one computes $H(V_0, V_1)$ and verifies that the result equals V , the root value.

[0024] Notice that a Merkle tree needs not to be a full binary tree. (For instance, if it is not, one can add dummy nodes to make it a full binary tree. Alternatively, instead of creating artificial nodes, it is possible to deem the value corresponding to any missing node a special, predetermined value, denoted by **EMPTY** in the sole figure.) Of course, if the tree is not sufficiently full, then authentication paths may become needlessly long.

[0025] One way to store certificate information in a Merkle tree includes associating pieces of the certificate information to leaves of the Merkle tree. For instance, the pieces of information are the values corresponding to the leaf nodes, while the one-way hash function determines the values of internal nodes of the tree, including the root. (As discussed elsewhere herein, however, it is possible to make certificate information correspond

to internal nodes, including the root.)

[0026] Each node of the Merkle tree that corresponds to a portion of the certificate information is deemed a "certificate node". A simple way to associate certificate information to certificate nodes includes having as many certificate nodes as there are certificates, and making information about an individual certificate be the value corresponding to an individual certificate node.

[0027] Instead of having the CA provide information directly to users, in a preferred embodiment, one or more intermediaries interact with most of the users. The one or more intermediaries obtain the certificate information from the CA. (Within an authenticated tree, we refer to a value belonging to the authentication path as *authenticating value*.) The CA may provide an intermediary with an entire *authenticated tree* indicating which certificates have been revoked. As illustrated below, an authenticated tree is a Merkle tree (storing certificate information) having a root value that is authenticated (e.g., digitally signed) by the CA. To prove to an end user that a given certificate is revoked, the intermediary provides the user with (a) the value of certificate node N of the authenticated tree, where N is the certificate node whose value indicates that the certificate in question has been revoked; (b) the authentication path of node N in the tree; and (c) the digital signature of the CA of the root value of the tree. (Recall that this signature may include date information and additional information).

[0028] The end user, upon receiving (a), (b), and (c) from the intermediary, verifies that the digital signature of the CA is valid, thus learning the true root value of the tree. Then, the user also verifies that the authentication path for node N is valid with respect to the root value, thus learning the true information about the certificate in question. (Here by "true" we mean deemed true by the CA.) The user also verifies that the date information, if any, of the CA's signature of the root value is the expected date. (For instance, if the CA includes in the date information not only the date in which the authenticated tree was constructed, but also the date by which the CA intends to update the authenticated tree, the user also verifies that he is dealing with the latest authenticated tree.)

[0029] Notice that if the user trusts the CA, the user needs not trust the intermediary. Indeed, if the intermediary wishes to provide the user with false information about the certificate in question, the intermediary needs to perform an extraordinary amount of computation. For instance, if the intermediary wishes to modify the value corresponding to node N , the intermediary would have to either change the root value and therefore be able to forge the digital signature of the CA, or be able to break the one-way hash function.

[0030] Notice that the amount of data that an honest intermediary provides to a user in order to allow the user to verify that the revocation information about a given certificate is authentic is not very large, especially when compared to a conventional CRL. Indeed, the bulk of

the data provided to the user consists of the authentication path of the node N. For instance, if there are 3,000 revoked certificates in the tree, then certificate information about these 3,000 certificates could be associated to the leaves of a Merkle tree of depth 12. Thus, there are at most 12 values in the authentication path of node N. If each value is 200 bits long (which is more than currently advocated for one-way hash functions), the total length of the authentication path is about 2400 bits. (more precisely, it is 11 times 200 bits for the values of internal nodes plus the length of the value of a leaf node. In fact, the authentication path of a certification leaf consists of one leaf value and eleven internal node values. As we shall see herein, however, even the length of leaf values can be made to be relatively small for authentication purposes.) This is much shorter than a CRL comprising the serial numbers and revocation dates of 3,000 revoked certificates. Of course, the new system also includes a digital signature of the root and date information, but so does the CRL system. Therefore, we succeed in a much more efficient way than a CRL system in having authentic information about certificates be provided to the end user by an intermediary without having to trust the intermediary.

[0031] In the example given above, an authenticated tree is used by the CA to provide the intermediary with data enabling the intermediary to prove to end users that a given certificate of the CA has been revoked. The CA may also use authenticated trees so as to enable the intermediary to prove much more general certificate information. For instance, as we have said, it should be possible for an intermediary to prove to end users not only the certification status of issued certificates, but also that (alleged) certificates have never been issued by the CA. For instance, this could be done as follows: Assume that, like in the PKI envisaged in the MITRE study, a serial number consists of a twenty-bit string. Then, each possible serial number can be put in correspondence with a leaf of a full binary tree of the depth twenty. Then, for each serial number, X, the CA creates an X-value such as "certificate number X has not yet been issued", or "certificate number X (has been issued and) is currently valid", or "certificate number X has been revoked (on date dx and for reason rx)", whichever is the case. Each X value is then associated to a leaf of its own. Once more, in this example, certificate nodes are leaf nodes. Then the CA associates values to all other nodes of the binary tree of level twenty using a one-way hash function H so as to form a Merkle tree. Then, the CA digitally signs the root value, thus obtaining an authenticated tree. An intermediary can use such a tree to efficiently prove to any end user the certification status of any possible serial number for a certificate, and thus the status of any possible certificate.

[0032] Note that it is possible to use the position of a node in an authenticated tree to convey information. For instance, in the above example, the CA may place information about certificate number X in leaf X so as to

avoid having to store the serial number in the leaf. (Indeed, a leaf of a binary tree of depth twenty can be associated to a unique twenty-bit string in a natural way. Consider the path from the root to the leaf. If the first node from the root is a left child, we write "0". If the next node is a right child, we write "1". Thus, upon reaching the leaf in question, we have written a twenty-bit string X that uniquely identifies the leaf.) Position information in the tree can be exploited using different ways of encoding positional information. For instance, if the serial numbers of certificates in the tree begin with number one, then certificate information for certificate serial number one can be placed in the first leaf (i.e., left most terminal leaf), certificate information about certificate serial number two can be placed in the second leaf, etc.

[0033] In addition, it is possible to put information about more than one certificate into a single node of the tree so that, for example, if information about two certificates were placed in each of the terminal leaf nodes, then information about certificate serial numbers one and two could be in the first terminal leaf node, information about certificate serial numbers three and four could be placed in the second terminal leaf node, etc.

[0034] An authority that issues and revokes certificates maps certificate information into values. This mapping may be as simple (e.g., the identity mapping) or complex as required by the nature of the certificate information. For example, if the certificate information consists simply of identifying whether a certificate is issued, revoked, or expired, then the mapping could be 00=issued, 01=revoked, (10=expired,) and 11=never issued. In this way, each possible value corresponds to different possible states for each certificate.

[0035] As can be seen in the above example, it is possible to store, in a certificate node, certificate information about more than one certificate. If the values for M certificates can be stored in each node, and there are 2^{20} possible certificates, then the Merkle tree may have $2^{20}/M$ leaf nodes. A proper encoding is used so that the value a certificate node can be understood without ambiguity. For instance, if we care about proving information of the type 00=issued, 01=revoked, 10=expired, and 11=never issued, and we want every value in the authenticated tree to be 200 bits long, then information for one hundred certificates may be stored in each certificate node. One knows a priori that two bits of certificate status about certificate number one through one hundred are to be found in the value of the first leaf. Further, one knows that the first two bits of that value correspond to certificate number one, that the second pair of bits in that value correspond to certificate number two, etc.

[0036] As discussed above, it is not necessary for the users to trust the intermediary since the intermediary provides certificate information to the users in a way that indicates that the CA has authenticated the certificate information. Note that an intermediary includes, but is not limited to, a directory, a distributor, a redistributor, a user, a CA, a database, a readable computer file, a read-

only computer file, an entity that has obtained information from another intermediary, or any combination of the above. Generally, an intermediary is an entity that provides certificate information authenticated by a CA.

[0037] A major advantage of conveying certificate information via authenticated trees consists of efficient updating. For instance, in a typical day, a CA may issue ten more certificates and revoke one previously issued certificate. In such a case, the CA may provide the intermediary with a new authenticated tree that contains the current certificate information about all certificates. More efficiently, however, the CA may provide the intermediary just with the new values of the certificate nodes, whose values have changed, together with the digital signature of the new root (and proper date information). In sum, therefore, the CA sends very little information to the intermediary. (Of course, if so wanted, the CA can send additional information. For instance, he may send not just the certified nodes that have changed, but values of nodes that have changed and the position of these nodes. So doing, he may simplify the work of the intermediary.)

[0038] The intermediary could use the new received values and the old values and the one-way hash function to compute the new Merkle tree, and thus, in particular, the new root value so that the intermediary can check the digital signature of the CA and verify the new authenticated tree. The intermediary is now ready to provide a user with updated certificate information.

[0039] The same efficiency can be obtained when the intermediary wishes to provide end users proofs of certificate information. For instance, a given user U may already be in possession of yesterday's entire authenticated tree. In this case, the intermediary can bring the user up to date by providing the user with just the values that have changed and the signature of the CA of the new root value. Of course, the user may not have yesterday's full authenticated tree, but he may have the full authenticated tree of, say, two weeks earlier. In this case, the intermediary may still provide the user with just the values that have changed with respect to two weeks ago along with the CA's signature of the new root value. This may still result in substantial savings.

[0040] The intermediary may know, by keeping track of what was sent to the user, which authenticated tree the user already knows. Alternatively, the user may signal to the intermediary which authenticated tree he already has and the intermediary may act accordingly. In general, the intermediary may omit sending to the user values that the user already knows. In addition, it is possible for the intermediary to obtain some of the authenticated tree information from an entity other than the CA, in which case the CA may only provide the intermediary with authenticated tree information that is not already known by the intermediary. Note that it is also possible for the intermediaries to obtain information about the tree from other sources, such as other intermediaries.

[0041] In another embodiment of the invention, the

CA may make use of two or more Merkle trees to convey certificate information. For instance, a first tree may contain information about issued but non-revoked certificates, and a second tree may contain information about revoked certificates. The root values of these trees (together with other information deemed proper) may be digitally signed by the CA either separately or together, thus making these Merkle tree authenticated trees.

[0042] Alternatively, it is possible to have a first authenticated tree containing information about all issued certificates and a second authenticated tree containing information about all revoked certificates. This embodiment is similar to the two tree system described above where one authenticated tree contains information about issued but non-revoked certificates and the second authenticated tree contains information about revoked certificates. However, in the case, knowing the issued certificates and revoked certificates makes it possible to determine or ascertain if a certificate is valid (i.e., by determining if the certificate is both issued and not revoked). It is also possible to construct an authenticated tree containing information about certificates that have not been issued. In that case, the authority can provide an intermediary with information for proving that one or more certificate serial numbers (or other appropriate certificate identifiers) do not correspond to any certificate that was issued as of a given date. Other possible combinations of authenticated trees containing certificate information are possible, as will become apparent from the following discussion.

[0043] Note that an authenticated tree may be constructed by the CA or by another entity, such as the intermediary, and that the other entity may simply present the root of the underlying Merkle tree to the CA for authentication. Also, it is possible for the CA to authenticate nodes, other than or in addition to the root, of a Merkle tree. In the case of a node other than the root of a Merkle tree being authenticated by the CA, an authentication path can be verified with respect to such an authenticated node rather than with respect to the root. It should be noted, however, that by authenticating (e.g., digitally signing) one or more nodes of a Merkle tree, the CA is actually generating authenticated (sub)trees that are stand-alone authenticated trees in their own right.

[0044] Note that the certificate information CI corresponding to a certificate node N may be one-way hashed prior to being associated to the node. That is, the value associated to node N is $H(CI)$ rather than CI. The process of passing from CI to $H(CI)$ is a possible mapping step performed by a CA. One advantage of such a mapping is that if node N is, say, a leaf node, then whenever the value of N is used within an authentication path of another node, it only contributes some, say, 200 bits to such a path no matter how long CI may be. (Indeed, CI may be a very long string indicating, among other things, the reasons that a certificate has been revoked.) On the other hand, when we wish to provide CI in an authenticated manner, we reveal CI in full.

(See, for example, Leighton et al.). Such a full value of CI is first mapped to $H(CI)$ and then this mapped value is proved to be the value of a genuine certificate node in an authenticated tree.

[0045] Note also that the intermediary may not directly send data to an end user, but cause the end user to receive data (e.g., by someone else, possibly at a signal of an intermediary, or by enabling the user to read some data file, or in any other manner).

[0046] Note that the scheme disclosed herein can be expanded beyond the simple type of tree disclosed in Merkle. An example of such expansions are given in U.S. patent no. 5,432,852 to Leighton et al, which is incorporated by reference herein. In particular, although an authenticated tree may have an underlying structure of a tree such as that disclosed in Merkle, the authenticated tree may in fact may have other interconnections and paths therethrough. Similarly, it is possible to have the value of a node depend on the position, P , of the node within the tree. For example, in such a case, the value of the node may be $H(VL, VR, P)$, where VL is the value of the node's left child and VR is the value of the node's right child. Including position information helps prevent attack by someone who attempts to find two values that hash to an actual node value of an authenticated tree. When the node information includes position information, then a prospective attacker would have to find values that hashed to the combination of the children information and the position information, which is much less likely. U.S. patent no. 5,432,852 to Leighton et al. discloses a similar idea used in connection with a signature scheme. Note also that the hashed value of a node could be a hash of the children concatenated with a position of the node within the tree or could be a hash of the concatenation of the children and the position.

[0047] In addition, one or more of the nodes containing the certificate information (i.e., the certificate nodes), are not all necessarily leaf nodes of the tree, but, instead, have children of their own. For instance, let N be an internal node whose left child has a value VL and whose right child has a value VR , and let CI be some certificate information that we wish to associate to node N . Then, the value associated with node N can be made equal to $VN = H(VL, VR, CI)$.

[0048] It should be realized the system disclosed herein facilitates batch processing. For example, rather than signing n pieces of certificate information separately, a CA may "Merkle" hash the quantities so as to obtain a single root-value, and then digitally signs the root-value (together with additional quantities if desired). This is a way to substitute n individual signatures (that can be expensive to obtain) with just one signature and $n-1$ hashing (which are not expensive at all). Notice that any signature scheme can be used for the root, including pen-written signatures, rather than digital ones.

[0049] Notice that a one-way hash function needs not to map every string to values having always B bits. For instance, it may map some strings to 160-bit values and

some other strings to 200-bit values. Also notice that an authenticated tree needs not to be constructed by means of a single one-way hash function. For instance, a first one-way hash function could be used for nodes at the first level, a second one-way hash function for nodes at the second level, etc. More generally, one could use a different one-way hash function for each position within the tree. Notice, however, that such a collection of one-way hash functions is a single one-way hash function for purposes of constructing an authenticated tree.

[0050] The foregoing has outlined some of the more pertinent objects and details of a preferred embodiment of the present invention. These objects and details should be construed to be merely illustrative of some of the more prominent features and applications of the invention. Many other beneficial results can be attained by applying the disclosed invention in a different manner or modifying the invention as will be described. Those skilled in the art will recognize that the invention can be practiced, with modification, in other and different certification methods and schemes within the spirit and scope of the invention. Also, note that the need for certification and certificate revocation extends beyond certifying public keys and could include certifying any information.

[0051] One of the preferred implementations of the various routines disclosed above is as a set of instructions in a code module resident in the random access memory of a computer. Until required by the computer, the set of instructions may be stored in another computer memory, for example, in a hard disk drive, or in a removable memory such as an optical disk (for eventual use in a CD ROM) or floppy disk (for eventual use in a floppy disk drive). In addition, although the various methods described are conveniently implemented in a general purpose computer selectively activated or reconfigured by software, one of ordinary skill in the art would also recognize that such methods may be carried out in hardware, in firmware, or in more specialized apparatus constructed to perform the required method steps.

45 Claims

1. A method for a intermediary to prove to an end user that certificate information is authenticated by an authority, comprising the steps of:

(a) obtaining at least a portion of an authenticated tree having certificate nodes corresponding to certificate values indicative of the information; and

(b) causing the end user to receive at least one of the following values: certificate values, authenticating values of certificate values, and one or more node values authenticated by an

authority.

2. A method according to claim 1, wherein at least one of the authenticated node values is the root value.
3. A method according to claim 1, wherein the user receives at least one of: a certificate value, a node value, and an authenticated node value.
4. A method according to claim 1, wherein the user uses at least a value that the intermediary caused the user to receive to verify the authenticity of the certificate information.
5. A method according to claim 4, wherein the end user verifies the authenticity of the certificate information via an authentication path of at least one certificate node.
6. A method according to claim 1, wherein the authenticated tree values and the authenticated node values change over time.
7. A method according to claim 6, wherein the intermediary causes the end user to receive values that have changed.
8. A method according to claim 6, wherein the intermediary omits causing the user to receive at least one value, already known to the user, belonging to an authentication path of a certificate node corresponding to certificate information that the user wants to verify.
9. A method according to claim 8, wherein the intermediary omits causing the user to receive a value corresponding to a node other than the root.
10. A method according to claim 6, wherein the user asks for the value of at least one certificate node and wherein the intermediary omits causing the user to receive at least one value in an authentication path of a certificate node already known to the user.
11. A method according to claim 1, wherein the intermediary sends at least one of: a certificate value, a node value, and an authenticated node value.
12. A method according to claim 11, wherein the intermediary omits sending at least one value already known to the user.
13. A method according to claim 12, wherein the user sends a signal indicating values that the user already knows.
14. A method according to claim 7, wherein the intermediary sends at least one of: a certificate value, a node value, and an authenticated node value.
15. A method according to claim 14, wherein the intermediary omits sending at least one value already known to the user.
16. A method according to claim 15, wherein the user sends a signal indicating values that the user already knows.
17. A method according to claim 1, wherein the intermediary can prove to the user that a certificate was never issued.
18. A method according to claim 1, wherein the intermediary can prove to the user that a alleged certificate was never issued.
19. A method according to claim 1, wherein the intermediary can prove to the end user that a given serial number does not correspond to any issued certificate of a given CA.
20. A method according to claim 1, wherein the certificates are public key certificates.
21. A method for providing authenticated information about a revoked certificate, comprising the steps of:
 - (a) providing an information value stored in an authentication tree, wherein the information value contains information about the revoked certificate;
 - (b) providing an authenticated root value of the authentication tree; and
 - (c) providing an authentication path for the information value.
22. A method according to claim 21, wherein the information value includes the serial numbers of the revoked certificate.
23. A method according to claim 22, wherein the information value includes a revocation date.
24. A method according to claim 21, wherein the root value is authenticated by a digital signature.
25. A method according to claim 24, wherein the digital signature is provided by a certification authority.
26. A method according to claim 24, wherein the digital signature is provided by a directory.
27. A method according to claim 21, wherein the information value is stored by a certification authority.
28. A method according to claim 21, wherein an entity

that stores the information value also digitally signs the root of the authentication tree.

29. A method according to claim 28, wherein the entity is a certification authority. 5
30. A method according to claim 28, wherein the entity is a directory.
31. A method according to claim 21, wherein the authentication path is provided by a directory. 10
32. A method according to claim 31, wherein the directory stores the information value and authenticates the root value. 15
33. A method according to claim 21, wherein the information about the revoked certificate is stored in a node of the authentication tree that depends upon the revoked certificate. 20

25

30

35

40

45

50

55

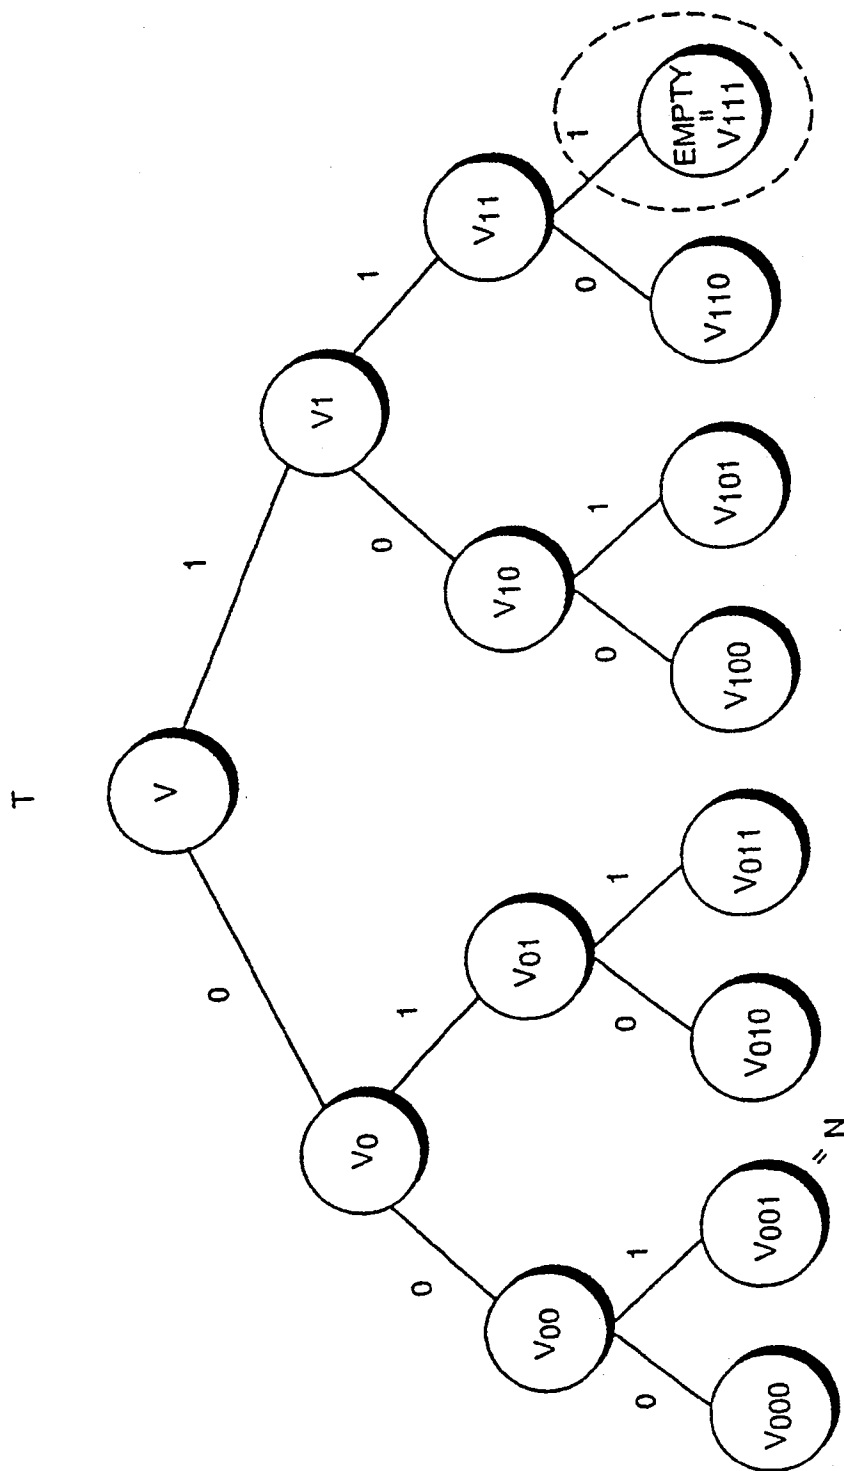



FIG. 1



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) **EP 1 164 746 A3**

(12)

EUROPEAN PATENT APPLICATION

(88) Date of publication A3:
30.10.2002 Bulletin 2002/44

(51) Int Cl.7: **H04L 9/32**

(43) Date of publication A2:
19.12.2001 Bulletin 2001/51

(21) Application number: **01119418.0**

(22) Date of filing: **01.11.1996**

(84) Designated Contracting States:
**AT BE CH DE DK ES FI FR GB GR IE IT LI LU MC
NL PT SE**

(71) Applicant: **Novomodo, Inc.**
Brookline, Massachusetts 02146 (US)

(30) Priority: **02.11.1995 US 6143 P**
11.10.1996 US 729619

(72) Inventor: **Novomodo, Inc.**
Brookline, Massachusetts 02146 (US)

(62) Document number(s) of the earlier application(s) in
accordance with Art. 76 EPC:
96937813.2 / 0 858 702

(74) Representative: **Bickel, Michael**
Westphal - Mussgnug & Partner
Patentanwälte
Mozartstrasse 8
80336 München (DE)

(54) **Tree-based certificate revocation system**

(57) A method for a intermediary to prove to an end user that certificate information is authenticated by an authority, comprising the steps of:

(a) obtaining at least a portion of an authenticated tree having certificate nodes corresponding to cer-

tificate values indicative of the information; and
(b) causing the end user to receive at least one of the following values: certificate values, authenticating values of certificate values, and one or more node values authenticated by an authority.




FIG. 1

EP 1 164 746 A3



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 01 11 9418

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
D,A	US 4 309 569 A (MERKLE RALPH C) 5 January 1982 (1982-01-05) * column 2, line 32 - column 3, line 2 * * column 3, line 50 - column 4, line 4 *	1,21	H04L9/32
A	SANTOSH CHOKHANI: "TOWARD A NATINAL PUBLIC KEY INFRASTRUCTURE" IEEE COMMUNICATIONS MAGAZINE, IEEE SERVICE CENTER, PISCATAWAY, N.J, US, vol. 32, no. 9, 1 September 1994 (1994-09-01), pages 70-74, XP000476557 ISSN: 0163-6804 * page 71, right-hand column, line 1 - page 72, right-hand column, line 12 * * page 73, left-hand column, last paragraph *	1,21	
			TECHNICAL FIELDS SEARCHED (Int.Cl.7)
			H04L
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 3 September 2002	Examiner Holper, G
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document		T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same parent family, corresponding document	

EPO FORM 1503 03/02 (P/C/M)

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 01 11 9418

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.
The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

03-09-2002

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
US 4309569	A	05-01-1982	NONE	

EPO FORM P0458

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82